

PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau



INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 13/18	A1	(11) International Publication Number: WO 95/15528
		(43) International Publication Date: 8 June 1995 (08.06.95)
(21) International Application Number: PCT/US94/13551		(81) Designated States: JP, KR, European patent (AT, BE, CH, DE, DK, ES, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).
(22) International Filing Date: 23 November 1994 (23.11.94)		
(30) Priority Data: 08/158,979 30 November 1993 (30.11.93) US		
(71) Applicant: VLSI TECHNOLOGY, INC. [US/US]; 1109 McKay Drive, San Jose, CA 95131 (US).		
(72) Inventors: RHODEN, William, Desi; 3412 East Suncrest Court, Phoenix, AZ 85044 (US). WU, Shih-Ho; 1360 W. Isabella Avenue, #2004, Mesa, AZ 85202 (US). NAKA-HARA, Mike; 13021 South 48th Street, #2044, Phoenix, AZ 85044 (US).		
(74) Agent: KREBS, Robert, E.; Burns, Doane, Swecker & Mathis, Washington and Prince Streets, P.O. Box 1404, Alexandria, VA 22313-1404 (US).		

Published

With international search report.


Before the expiration of the time limit for amending the claims and to be republished in the event of the receipt of amendments.

(54) Title: A REALLOCATABLE MEMORY SUBSYSTEM ENABLING TRANSPARENT TRANSFER OF MEMORY FUNCTION DURING UPGRADE

(57) Abstract

A low-cost, moderate performance small computer system is provided by allowing a single sharable block of memory to be independently accessible as graphics or main store memory. Allocation of the memory is selected programmably, eliminating the need to have the maximum memory size for each block simultaneously. Performance penalties are minimized by dynamically allocating the memory bandwidth on demand rather than through fixed time slices. A reallocatable memory subsystem enables transparent transfer of memory function of a lower-performance memory such as

DRAM to occur in conjunction with a memory upgrade to a higher-performance memory such as VRAM, for example. More particularly, an apparatus for use in a computing machine including a CPU and a first memory includes memory slots, configuration circuitry for enabling allocation of a second memory as display memory and the first memory as main memory responsive to a second memory having a different performance level than the first memory being added to the memory slots; and circuitry for allowing substantially independent access to the first memory and the second memory. Circuitry for allowing substantially independent access may include a memory controller including arbitration circuitry for arbitrating among a plurality of requests for access to the first and second memories, a first data path connected to the arbitration circuitry and including a first buffer store for facilitating exchange of data with the first memory, a second data path connected to the arbitration circuitry and including a second buffer store for facilitating exchange of data with the second memory, and control circuitry connected to the configuration circuitry and responsive to one or more signals applied to the apparatus, the signals including address, data and control signals, for causing at least some of the data signals to be applied to only one of the first and second data paths.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AT	Austria	GB	United Kingdom	MR	Mauritania
AU	Australia	GE	Georgia	MW	Malawi
BB	Barbados	GN	Guinea	NE	Niger
BE	Belgium	GR	Greece	NL	Netherlands
BF	Burkina Faso	HU	Hungary	NO	Norway
BG	Bulgaria	IE	Ireland	NZ	New Zealand
BJ	Benin	IT	Italy	PL	Poland
BR	Brazil	JP	Japan	PT	Portugal
BY	Belarus	KE	Kenya	RO	Romania
CA	Canada	KG	Kyrgyzstan	RU	Russian Federation
CF	Central African Republic	KP	Democratic People's Republic of Korea	SD	Sudan
CG	Congo	KR	Republic of Korea	SE	Sweden
CH	Switzerland	KZ	Kazakhstan	SI	Slovenia
CI	Côte d'Ivoire	LI	Liechtenstein	SK	Slovakia
CM	Cameroon	LK	Sri Lanka	SN	Senegal
CN	China	LU	Luxembourg	TD	Chad
CS	Czechoslovakia	LV	Latvia	TG	Togo
CZ	Czech Republic	MC	Monaco	TJ	Tajikistan
DE	Germany	MD	Republic of Moldova	TT	Trinidad and Tobago
DK	Denmark	MG	Madagascar	UA	Ukraine
ES	Spain	ML	Mali	US	United States of America
FI	Finland	MN	Mongolia	UZ	Uzbekistan
FR	France			VN	Viet Nam
GA	Gabon				

**A REALLOCATABLE MEMORY SUBSYSTEM ENABLING
TRANSPARENT TRANSFER OF MEMORY FUNCTION
DURING UPGRADE**

BACKGROUND OF THE INVENTION

5 Field of the Invention:

The present invention relates generally to computer architecture, and more particularly, to memory-sharing architectures which include graphics capabilities.

State of the Art:

10 As the density of solid state memories increases, oversized memories are being wastefully used for purposes which optimally require specialized memory configurations (e.g., graphics refresh). One reason for this is that manufacturers attempt to produce memory sizes which will achieve a broad range of applicability and a high volume of production. The more popular, and thus more cost-effective memories, tend to be fabricated with square aspect ratios or with
15 tall, thin aspect ratios (i.e., a large number of fixed length words) that are not readily suited to specialized uses.

20 Although uses which can exploit memories with these popular aspect ratios can be implemented in a relatively cost-effective manner, specialized uses which cannot exploit these aspect ratios can be proportionately more expensive to implement. The expense associated with implementing specialized uses assumes one of two forms: (1) the increased cost associated with purchasing a memory which does not conform to a readily available and widely used memory configuration; or (2) the increased cost associated with purchasing a readily
25 available memory which is much larger than needed specialized use (e.g., a relatively square memory which must be tall enough to obtain a desired width,

-2-

even though only a relatively small number of rows in the memory are needed for the purpose at hand).

5 The foregoing memory capacity problem is typically referred to as the memory granularity problem: expensive chips can be purchased and used efficiently or inexpensive memory chips can be purchased and used inefficiently. This problem is especially significant in computer systems which implement graphics functions, since these systems typically include a dedicated, high speed display memory. Specialized display memories are usually required because typically refresh function for the graphics display (e.g., for a 1280 x 1024 display) consumes virtually all of the available bandwidth of a typical dynamic random access memory (DRAM).

15 To update a video line on a high resolution graphics display, a graphics refresh optimally requires a memory having a short, wide aspect ratio. Display memories used as frame buffers for high resolution graphics displays have therefore become an increasingly larger fraction of a system's overall cost due to the foregoing memory problem. For display memories, even a two megabit memory can be unnecessarily large, such that it cannot be effectively used. An exemplary display memory for a current high-end display of 1280 x 1024 pixels requires just over one megabyte of memory. Thus, almost one-half of the display memory remains unused.

20 For example, Figure 1 illustrates a typical computer system 100 which includes graphics capabilities. The Figure 1 computer system includes a central processing unit (CPU) 102, a graphics controller 104 and a system controller 106 all connected to a common bus 108 having a data portion 110 and an address portion 112.

25 The graphics controller 104 is connected to display memory 114 (e.g., random access memory, or RAM) by a memory bus having a memory address

-3-

bus 116 and a memory data bus 118. RAMDAC 120 performs digital-to-analog conversion (DAC) of signals (e.g., analog RGB color signals) used to drive a graphics display.

5 The system controller is connected to system memory 122 by a separate memory address bus 124. A memory data bus 126 is connected directly between the common data bus 108 and the system memory. The system memory can also include a separate cache memory 128 connected to the common bus to provide a relatively high-speed portion for the system memory.

10 The graphics controller 104 mediates access of the CPU 102 to the display memory 114. For system memory transfers not involving direct memory access (DMA), the system controller 106 mediates access of the CPU 102 to system memory 122, and can include a cache controller for mediating CPU access to the cache memory 128.

15 However, the Figure 1 configuration suffers significant drawbacks, including the granularity problem discussed above. The display memory 114 is limited to use in connection with the graphics controller and cannot be used for general system needs. Further, because separate memories are used for the main system and for the graphics memory, a higher number of pin counts render integration of the Figure 1 computer system difficult. The use of separate
20 controllers and memories for the main system and the graphics also results in significant duplication of bus interfaces, memory control and so forth, thus leading to increased cost. For example, the maximum memory required to handle worst case requirements for each of the system memory and the graphics memory must be separately satisfied, even though the computer system will
25 likely never run an application that would require the maximum amount of graphics and main store memory simultaneously. In addition, transfers between the main memory and the graphics require that either the CPU or a DMA controller intervene, thus blocking use of the system bus.

-4-

Attempts have been made to alleviate the foregoing drawbacks of the Figure 1 system by integrating system memory with display memory. However, these attempts have reduced duplication of control features at the expense of system performance. These attempts have not adequately addressed the granularity problem.

Some attempts have been made, particularly in the area of portable and laptop systems, to unify display memory and system memory. For example, one approach to integrated display memory and system memory is illustrated in Figure 2. However, approaches such as that illustrated in Figure 2 suffer significant drawbacks. For example, refreshing of the display via the graphics controller requires that cycles be stolen from the main memory, rendering performance unpredictable. Further, these approaches use a time-sliced arbitration mode for allocating specific time slots among the system controller and the graphics controller, such that overall system performance is further degraded.

In other words, overall performance of the Figure 2 system is limited by the bandwidth of the single memory block, and the high demands of graphics refresh function alone introduce significant performance degradation. The allocation of memory bandwidth between display access and system access using fixed time-slots only adds to performance degradation. Because the time slots must be capable of handling the worst case requirements for each of the system memory and display memory subsystems, the worst possible memory allocation is forced to be the normal case.

Examples of computers using time-slice access to an integrated memory are the Commodore and the Amiga. The Apple II computer also used a single memory for system and display purposes. In addition, the recently-released Polar™ chip set of the present assignee, for portable and laptop systems, makes provision for integrated memory.

-5-

A different approach is described in a document entitled "64200 (Wingine™) High Performance 'Windows™ Engine'", available from Chips and Technologies, Inc. In one respect, Wingine is similar to the conventional computer architecture of Figure 1 but with the addition of a separate path that enables the system controller to perform write operations to graphics memory. The graphics controller, meanwhile, performs screen refresh only. In another respect, Wingine may be viewed as a variation on previous integrated-memory architectures. Part of system memory is replaced with VRAM, thereby eliminating the bandwidth contention problem using a more expensive memory (VRAM is typically at least twice as expensive as DRAM). In the Wingine implementation, VRAM is not shared but is dedicated for use as graphics memory. Similarly, one version of the Alpha microprocessor sold by Digital Equipment Corporation reportedly has on board a memory controller that allows VRAM to be used to alleviate the bandwidth contention problem. The CPU performs a role analogous to that of a graphics controller, viewing the VRAM frame buffer as a special section of system RAM. As with Wingine, the VRAM is not shared.

Thus, traditional computer architectures can not efficiently integrate a single memory to accommodate the two different functions of display memory and system memory without significantly degrading system performance. What is needed, then, is a new computer architecture that allows display memory and system memory to be integrated while still achieving high system performance. Such an architecture should, desirably, allow for memory expansion and use with cache memory. Further, any such system should provide an upgrade path to existing and planned high performance memory chips, including VRAM, synchronous DRAM (SDRAM) and extended data out DRAM (EDODRAM).

-6-

SUMMARY OF THE INVENTION

The present invention, generally speaking, provides a low-cost, moderate performance small computer system by allowing a single sharable block of memory to be independently accessible as graphics or main store memory.

5 Allocation of the memory is selected programmably, eliminating the need to have the maximum memory size for each block simultaneously. Performance penalties are minimized by dynamically allocating the memory bandwidth on demand rather than through fixed time slices.

10 A reallocatable memory subsystem enables transparent transfer of memory function of a lower-performance memory such as DRAM to occur in conjunction with a memory upgrade to a higher-performance memory such as VRAM, for example. More particularly, an apparatus for use in a computing machine including a CPU and a first memory includes memory slots, configuration
15 circuitry for enabling allocation of a second memory as display memory and the first memory as main memory responsive to a second memory having a different performance level than the first memory being added to the memory slots; and means for allowing substantially independent access to the first memory and the second memory. Circuitry for allowing substantially independent access may
20 include a memory controller including arbitration circuitry for arbitrating among a plurality of requests for access to the first and second memories, a first data path connected to the arbitration circuitry and including a first buffer store for facilitating exchange of data with the first memory, a second data path connected to the arbitration circuitry and including a second buffer store for facilitating
25 exchange of data with the second memory, and control circuitry connected to the configuration circuitry and responsive to one or more signals applied to the apparatus, the signals including address, data and control signals, for causing at least some of the data signals to be applied to only one of the first and second data paths.

-7-

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention can be further understood with reference to the following description and the appended drawings, wherein like elements are provided with the same reference numerals. In the drawings:

5 Figure 1 is a system block diagram of a conventional computer system;

Figure 2 is a block diagram of another conventional computer system;

Figure 3 is a system block diagram of a base computer system in accordance with an exemplary embodiment of the present invention;

10 Figure 4 is a more detailed block diagram of the graphics controller of Figure 3;

Figure 5 is a more detailed block diagram of the bus interface of Figure 3;

Figure 6 is a more detailed diagram of the bus status and configuration registers and decode block of Figure 5;

15 Figure 7 is a block diagram illustrating a remapping of memory in accordance with an exemplary embodiment of the present invention; and

20 Figure 8 is a system block diagram showing a system configuration in which the addition of VRAM has resulted in remapping of display memory to VRAM and reallocation of previous display memory to serve as system memory, in accordance with an exemplary embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

25 Figure 3 illustrates an exemplary embodiment of an apparatus for processing data in accordance with the present invention. The Figure 3 apparatus, generally labeled 300, can be a computer system which includes a main CPU 302. The main CPU 302 can, for example, be any available microprocessor, such as any standard 486-based processor.

-8-

5 The Figure 3 apparatus includes a means for storing data, generally represented as a memory 304. In accordance with the present invention, the data storing means 304 includes a system memory portion (e.g., random access memory, or RAM) and a display memory portion (e.g., RAM) addressed via common address lines 306 labeled MA. The display (e.g., graphics) memory portion can include an address space from an address 0 to an address (B-1) for a data storing means having B bytes. Further, the display memory portion and the system memory portion read and write data via common memory data lines 308 labeled MD.

10 The Figure 3 apparatus includes means for controlling a display operation of the Figure 3 system independently of the system controller. The display controlling means is generally represented as a display (e.g., graphics) controller 400. The graphics controller 400 is connected to the CPU 302 via CPU address lines 310 and CPU data lines 312 of a main CPU bus 314. The graphics
15 controller 400 controls access to the graphics memory portion of the data storing means.

The Figure 3 computer system further includes means for controlling access to the system memory portion of the data storing means 304. The means for controlling access to the system memory portion is generally represented as a
20 system controller 316 which is interfaced to the CPU 302 and the graphics controller 400 via the main CPU bus 314. Although the graphics controller and the system controller are indicated as separate blocks, in a physical implementation, they may reside on the same integrated circuit chip or on separate chips.

25 The signal lines 318, 322 and 324 permit the Figure 3 computer system to provide cache support for the system memory via the graphics controller 400, where the cache controller is included within the system controller. In accordance with exemplary embodiments, a cache memory 326 can be included

-9-

for this purpose. Memory reads and writes can be performed to the data storing means in both burst and non-burst modes.

5 Generally speaking, the signal line 322 labeled DRAM# indicates to the graphics controller when an addressable location exists within the shared memory and is not in the L2 cache. The signal line 324 labeled ERDY# is an early ready signal from the graphics controller to the system controller to verify that valid data has been read from the shared memory and will be valid for reading by the CPU in a predetermined time.

10 More particularly, typical personal computer systems feature an on-chip level-one (L1) cache of, for example, 8 kilo bytes within the CPU. Any external cache therefore functions as a level-two (L2) cache; i.e., data sought by the CPU is first sought in the L1 cache, then sought in the L2 cache, if necessary, and then sought in system memory if the data has not been found. In the conventional computer architecture of Figure 1, since system memory is located
15 in a single system memory 122, a cache controller included within the system controller 106 can function independently of the graphics controller 104.

In the system of Figure 3, on the other hand, system memory is located in the shared data storing means 304. However, in accordance with exemplary embodiments, existing cache control capabilities of the system controller 316 can
20 still be used by establishing communication between the graphics controller 400 and the system controller 316. Further, in the system of Figure 3, system memory is located in both the data storing means represented by memory 304, and an optional expansion memory 328. A failure to detect data in the L2 cache may therefore result in the data being found in the shared memory or in
25 expansion memory. Again, communication between the graphics controller 400 and the system controller 316 can handle this situation.

-10-

Figure 3 illustrates the manner in which efficient L2 cache memory support is provided for a system wherein a system controller 316 has an integrated L2 cache controller and a graphics controller, and a shared memory system. L2 cache support is provided for all system memory, regardless of the controller to which it is connected. Such support requires coordination between the system controller (with its integrated L2 cache controller) and the graphics controller.

In a 486-like or VL-Bus-based personal computer, L2 cache support may be provided using the existing backoff (i.e., BOFF#) CPU bus signal and the two new signals referred to herein as the DRAM# and ERDY# signals. DRAM# is driven by the system controller and ERDY# is driven by the graphics controller.

The system controller 316 monitors memory cycles and notifies the graphics controller when to ignore a particular memory cycle by deasserting the DRAM# on the signal line 322 at a predetermined time in the memory cycle. A system controller instructs the graphics controller to ignore a particular memory cycle when the addressable location is to a location other than the graphics portion of the data storing means (e.g., if the addressable location is to an ISA or PCI bus of the system, or if it's a location within the cache, or in another separate memory and so forth).

The graphics controller 400 also monitors memory cycles and begins a memory cycle when an addressable location is within the range of addressable locations for which the graphics controller is enabled to respond. In operation, the graphics controller tests the DRAM# on the signal line 322 at a predetermined time to determine whether it should respond to a current memory cycle. If the DRAM# signal on the signal line 322 has been deasserted by the system controller (i.e., false) the graphics controller 400 aborts the current memory cycle.

-11-

On the contrary, if the DRAM# on the signal line 322 has been asserted by the system controller (i.e., tests true), the memory cycle continues and the graphics controller 400 asserts the signal ERDY# on the signal line 324 to indicate to the system controller that the graphics controller is ready to read data. In this sense, the ERDY# signal represents an early ready signal which occurs a fixed number of clock cycles before data which is to be read becomes valid. In this instance, the cache controller 320 integrated within the system controller 316 senses the ERDY# signal on signal line 322 and initiates a writing of data into the cache 326.

10 The graphics controller can also be programmed to drive ERDY# at the end of a memory read cycle to signal to the system controller if a parity error occurred during the read.

Write-backs, for read-miss-dirty cycles and the like, are also supported using the BOFF# CPU bus signal. When write-back is required in response to a read request, the system controller asserts BOFF# (backoff), causing the CPU to abort the read cycle. Meanwhile, the graphics controller will have already started a memory read if the real address was within its address space.

15 The graphics controller also monitors BOFF# and, when it is asserted, is alerted that the read has been aborted. If the write-back is to memory outside the graphics controller's address space, the graphics controller may allow the read to continue, assuming that by the time the read has completed, the write-back may also be done, reducing latency time. The write-back may also be to memory in the graphics controller's address space. In this case, the system controller keeps BOFF# asserted and "masters" the write-back on the CPU bus by driving the bus just as the CPU would do if it were initiating the write. After the write-back has been completed, BOFF# is deasserted, and the CPU restarts the read operation.

-12-

This approach can be extended to provide L2 cache support for memory on other devices connected to the CPU bus. ERDY# may be driven by multiple sources in a "open-drain" configuration. Multiple DRAM# lines can be used or encoded together to signal to multiple devices.

5 In accordance with exemplary embodiments, the graphics controller 400 can include means for reallocating addressable locations of the data storing means 304 as display memory which is accessible by the graphics controller 400, or as system memory which is independently accessible by the system controller 316. Further, the exemplary graphics controller 400 can include means for
10 dynamically controlling access of the system controller and the display controlling means to the display memory portion and the system memory portion, respectively. The reallocating means and access controlling means are generally represented as block 500, included within the graphics controller 400.

15 The Figure 3 computer system can provide significant advantages. For example, the Figure 3 system represents a scalable architecture which can be configured for various price/performance alternatives. The Figure 3 system represents a relatively low-cost system which includes a single bank of shared memory (represented by the data storing means 304) which can be concurrently used, and dynamically reconfigured for both graphics and system functions.
20 Unlike previous shared memory systems, the allocation of memory bandwidth between display access and system access is not fixed; rather, memory bandwidth is dynamically allocated on demand between display access and system access.

25 Exemplary embodiments of the present invention, such as that illustrated in Figure 3, can achieve enhanced performance by adding a second bank of memory represented by the expansion memory means 328. In accordance with the exemplary embodiment wherein expansion memory is used, B bytes of memory in the shared memory can be allocated to system use, with an address space from address locations zero through address (B-1). The expansion memory

-13-

can be considered to contain E bytes of expansion system memory (e.g., RAM). In an exemplary embodiment, the E bytes can be addressed beginning with starting address B and ending with address $(E + B - 1)$.

5 In such an alternate embodiment, the data storing means 304 can continue to be shared between the graphics controller and the system controller. However, in accordance with alternate embodiments, a relatively high level of performance can be achieved by dedicating all of the data storing means 304 to graphics, reserving only the relatively fast portion of the data storing means or the expansion memory means for system use.

10 By the add on of expansion memory via an independent, separately controlled memory bus, system performance can be further enhanced, while using the same cache controller integrated in the system controller. With the addition of a simple memory interface block, concurrent accesses can occur to both the data storing means 304 and the expansion memory means 328. In this
15 case, performance can be further improved. For example, the possibility of parallel main memory accesses to two possible memory paths can result in increased performance by effectively overlapping accesses.

Thus, exemplary embodiments of the present invention provide significant advantages. By providing a single sharable block of memory that is
20 independently accessible as graphics memory or as main store memory, improved performance at relatively low-cost can be realized. By rendering allocation of the shared memory programmably selectable, any need to have maximum memory size for each of the independent graphics and main memory functions can be eliminated. Further, memory bandwidth can be dynamically allocated on
25 demand rather than via fixed time slices, further improving performance.

Referring to Figure 4, the graphics controller 400 interfaces to the CPU bus 314 via the reallocating means represented as bus interface 402. The

-14-

graphics controller interfaces to the data storing means 304 via the access controlling means, represented as a memory interface 408.

Commands and data from the Figure 3 CPU 302 are distributed to various logic blocks of the graphics controller 400 on two main buses represented by a display access bus 405 and a system access bus 407, indicated by thick, heavy lines in Figure 4. The system access bus 407 is connected to the memory interface 408.

The display access bus 405 is connected to various graphics controller logic blocks which are responsive to commands or programming instructions from the CPU. These logic blocks include a CRT controller (CRTC) 404, a sequencer (SEQ) 410, a RAMDAC interface 412, a clock synthesizer interface 418, an attribute controller (ATT) 422, a hardware cursor (HWC) 428, a graphics accelerator (Accel) 414 and pixel logic 416. In other implementations, other logic blocks may be included or ones of the foregoing logic block may not be included.

The CRTC 404 provides vertical and horizontal sync signals to a raster-scan CRT display. The sequencer 410 provides basic timing control for the CRTC 404 and the attribute controller 422. The RAMDAC interface 412 provides for programming of a RAMDAC (i.e., external or integrated) such as the RAMDAC of Figure 1. The RAMDAC is a combination random access memory and digital-to-analog converter that functions as a color palette which drives the CRT. The RAMDAC 120 in Figure 1 can be a look-up table used to convert the data associated with a pixel in the display memory into a color (e.g., RGB analog output).

The attribute controller 422 provides processing for alphanumeric and graphics modes. The hardware cursor 428 provides for display of any of a number of user-definable cursors. The accelerator 414 and pixel logic 416 assist

-15-

the host CPU in graphics-related operations. The pixel logic 416 of Figure 4 may also function as a pixel cache.

5 The clock synthesizer interface 418 provides for programming of a programmable clock synthesizer. Operation of the clock synthesizer interface, along with the other various graphics logic blocks in Figure 3, is well-known to one of ordinary skill in the art.

10 The memory interface 408, which functions as the access controlling means, arbitrates memory access between a number of different entities: the system access bus 407, the pixel logic 416, the display refresh logic 426, and the hardware cursor 428. Priority between these entities can vary according to system activity and the degree to which various buffers are full or empty. The priority scheme takes into account whether a particular access relates to a "mission-critical" function, so as to prevent such functions from being disrupted. For example, display refresh can be classified as a mission-critical function.

15 The exemplary Figure 3 system allocates a portion of the graphics controller's memory to the CPU for system use such that a single shared memory can be used to concurrently implement display functions and system memory functions. In accordance with alternate embodiments of the present invention, latency times for both graphics and system cycles can be further improved by providing separate queues for graphics and system accesses, with the separate
20 queues being serviced in parallel, independently of each other.

25 More particularly, Figure 5 shows the reallocating means represented by the bus interface 500 of Figure 4 in greater detail. As illustrated in Figure 5, a bus state machine 502 connects to the CPU bus and executes bus cycles involving the graphics controller. Commands or data from the CPU are latched in a command latch 504. The command latch is connected to both a graphics queue 506 and a system queue 508. The graphics queue 506 establishes bi-directional

-16-

operation using two separate, uni-directional queues: one queue that stores commands from the CPU and outputs them from the bus interface for use by the graphics controller, and one queue that stores data of the graphics controller and outputs it to the CPU. Likewise, the system queue 508 is a bi-directional queue composed of two unidirectional queues. The output buses of the graphics queue and the system queue are therefore bi-directional and are connected to an output latch 510 in order to drive data from the graphics controller to the CPU.

Separate memory and input/output (I/O) address ranges are defined for each queue such that the graphics and system queues are independently accessible. The graphics queue 506 and the system queue 508 are controlled by a graphics queue state machine 512 and a system queue state machine 514, respectively. These state machines are in turn controlled by the bus state machine 502.

A bus status/configuration registers/address decode block 520 is connected to the bus state machine 502. Further, block 520 is connected with an output multiplexer 516 of the output latch, and an output multiplexer ("mux") 518 of the command latch.

Bus status registers of block 520 contain information regarding the state of the graphics controller and the amount of available space in the graphics and system queues. The bus status registers may be read directly through the output mux 516 without putting a read command into either queue. Configuration registers of block 520 are written to from the bus state machine 502 and are used to select modes of operation in addition to those provided in a typical video graphics array (VGA) implementation.

In accordance with exemplary embodiments, programming flexibility can be improved by providing remapping registers which allow the CPU to reallocate the addresses to which the graphics controller responds. Address decoding is

-17-

programmable, such that the graphics controller responds to a CPU command if the command is to an address within the graphics controller's designated address space.

5 Outside the bus interface 402 of Figure 4, the graphics controller assumes that registers and memory are always at fixed addresses. Within the bus interface, address decode logic included in block 520 allows a register/memory location to be reallocated (i.e., remapped) from an original address to a new address more suitable to the CPU. This address decode logic therefore maps the new CPU address back to its original address.

10 An exemplary sequence would be as follows. The CPU issues a read command of a particular address. The graphics controller's address decode logic included in block 520 determines that the address is within the graphics controller's range, but that the desired register/memory location has been remapped from its original address to a new address more suitable to the CPU.
15 In this case, the address decode logic in block 520 maps the CPU address back to the original address and latches that address into the appropriate queue via the mux 518. Below the queues 506 and 508, registers and memory are always at fixed addresses, simplifying decoding of the graphics and system queue buses. In addition to the graphics queue 506 and the system queue 508, a separate latch
20 (one-stage queue) 522 can be provided for the hardware cursor.

Referring to Figure 6, the bus status/ configuration registers/address decode block 520 of Figure 5 is illustrated in greater detail. As shown in Figure 6, the block 520 includes address decode logic 602, configuration registers 604 and status registers 606. The address decode logic 602 examines the CPU
25 control lines that define whether the command is to memory or I/O and is a read or a write operation. The address decode logic 602 further compares the CPU address on the address bus to addresses programmed for various logic groups. If a match is found, the appropriate select line is asserted. Separate lines out of the

-18-

address decode logic signal if the CPU address is within the address space of one of the following exemplary groups: VGA mode I/O, VGA mode frame buffer, Windows mode registers, Windows mode frame buffer, system memory, configuration registers, or the status registers address space (which is within the configuration registers address space).

The configuration registers 604 are initialized to some pre-determined value at power-on reset. The configuration registers remap some of the address spaces within the graphics controller. This remapping allows software to access particular register or logic at a different address than to which it was initialized. Additional capability can be added to inhibit the graphics controller from responding to accesses of particular logic or memory. This may be done in various ways, for example, explicitly via enable/disable bits in a register and implicitly by programming the low and high address boundaries for a group to be the same. The configuration registers can be read by the CPU via a port 608.

The status registers 606 are read only. They contain information such as queue status (how full the queues are), what the accelerator is doing, what errors have occurred, and so forth. Certain bits of the status registers may be cleared by being read. The CPU reads the status registers directly without having to go through the graphics or system queues.

Figure 7 illustrates a reallocation of addressable locations in memory when the expansion memory means 328 of Figure 3 is used. The reallocation of Figure 7 ensures that addressable locations of any expansion memory are added to the bottom of available system memory. This ensures that expansion memory will always be accessed first by the CPU to accommodate system upgrades to high-speed memory.

The present reallocatable memory subsystem enables transparent transfer of DRAM function to occur in conjunction with a memory upgrade to VRAM,

-19-

for example. Existing systems typically allow upgrade of DRAM to faster DRAM or of DRAM to VRAM. Such systems waste any non-removable (i.e., motherboard-soldered) DRAM in the system as the new memory replaces the old. This waste is unavoidable, because the function of DRAM in such systems is fixed and cannot be transferred when the upgrade takes place.

Memory utilization is maximized (and total system cost minimized) by enabling transfer of DRAM function from graphics to main memory when a VRAM upgrade is added to the system.

Referring to Figure 8, the illustrated system is substantially identical to that of Figure 3 with the exception that a VRAM 805 has been added to the system and is connected to the graphics controller 400 by the memory address bus 306 and the memory data bus 308. When the more expensive VRAM is added to achieve higher performance, the function of display memory is migrated from the DRAM 804 (used in a basic system configuration as shared memory to provide adequate system performance at low cost) to the VRAM 805. The VRAM therefore takes over the function originally provided by the DRAM. The DRAM is transferred to a new function (e.g., system memory) and is still available to the system. No memory is wasted or lost as normally is the case. Graphics (and system) performance is significantly improved because of the additional bandwidth and feature set available from the VRAM 805. In particular, the VRAM has a serial port 807 that connects directly to the RAMDAC 809. The normal display refresh path described in conjunction with Figure 4 is therefore inactivated.

In a preferred embodiment, two bank select lines are dedicated to VRAM. At power-on, BIOS checks to see what memory is present in the system. If a memory device responds to these dedicated bank select signals, it is assumed to be VRAM. BIOS or a device driver then configures the graphics controller accordingly. In a preferred embodiment, 1M of VRAM or 2M of VRAM may

-20-

be added to the system. An analogous remapping as that shown in Figure 7 is performed to map display memory to the VRAM 805 and to map a portion of the system memory address space to the DRAM 804.

5 The memory interface 408 of Figure 4 distinguishes between DRAM and VRAM accesses to execute the correct physical cycles for the memory involved in a particular transfer.

Although an upgrade from DRAM to VRAM has been described in conjunction with an exemplary embodiment, the same mechanism may be used to upgrade to any of a number of existing and planned high performance memory chips, including VRAM, SDRAM and EDODRAM.

10

In summary, by integrating graphics memory and system memory, the present architecture allows system cost to be significantly reduced while providing an upgrade path. By enabling transparent transfer of memory function during upgrade, no memory is wasted or lost as is normally the case. High system performance is furthered by providing a bus interface with separate graphics and system paths.

15

It will be appreciated by those skilled in the art that the present invention can be embodied in other specific forms without departing from the spirit or essential characteristics thereof. The presently disclosed embodiments are therefore considered in all respects to be illustrative and not restricted. The scope of the invention is indicated by the appended claims rather than the foregoing description and all changes that come within the meaning and range and equivalence thereof are intended to be embraced therein.

20

-21-

WHAT IS CLAIMED IS:

1. For use in a computing apparatus including a CPU and a first memory, apparatus comprising:

memory slots;

5 configuration means for enabling allocation of a second memory as display memory and said first memory as main memory responsive to a second memory having a different performance level than said first memory being added to said memory slots; and

10 means for allowing substantially independent access to said first memory and said second memory.

2. The apparatus of Claim 1 further comprising:

memory controller means for controlling access to said first and second memories, including arbitration means for arbitrating among a plurality of requests for access to said first and second memories;

15 first data path means connected to said arbitration means and including first buffer storage means for facilitating exchange of data with said first memory;

20 second data path means connected to said arbitration means and including second buffer storage means for facilitating exchange of data with said second memory; and

control means connected to said configuration means and responsive to one or more signals applied to said apparatus, said signals including address, data and control signals, for causing at least some of said data signals to be applied to only one of said first and second data path means.

25 3. The apparatus of Claim 2 wherein said control means comprises a bus controller.

-22-

4. The apparatus of Claim 3 wherein said first buffer storage means comprises a system queue.

5. The apparatus of Claim 4 wherein said first buffer storage means further comprises a system queue controller connected to said system queue and to said bus controller.

5

1 / 6




FIG. 1
(PRIOR ART)




FIG. 2
(PRIOR ART)






FIG. 3

3 / 6

**FIG. 4**

**FIG. 5**

5 / 6






FIG._6

6 / 6

**FIG. 7****FIG. 8**

INTERNATIONAL SEARCH REPORT

Int. National Application No.

PCT/US 94/13551

A. CLASSIFICATION OF SUBJECT MATTER
IPC 6 G06F13/18

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 6 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X A	US,A,4 277 836 (KAWAKAMI) 7 July 1981 see column 1, line 15 - column 2, line 2 see column 2, line 25 - column 3, line 46 see abstract; claim 1; figures 1-3 ---	1-3 4,5
X A	EP,A,0 217 479 (NEC CORPORATION) 8 April 1987 see page 3 - page 5 see abstract; claim 1; figures 1-2 ---	1,2 3-5
A	US,A,4 028 663 (R.D. ROYER ET AL.) 7 June 1977 see column 1, line 64 - column 2, line 68 see column 3, line 46 - column 5, line 42 see abstract; claims 1-3; figures 1-3 -----	1-5

☐ Further documents are listed in the continuation of box C.☒ Patent family members are listed in annex.

* Special categories of cited documents :

- "A" document defining the general state of the art which is not considered to be of particular relevance
"E" earlier document but published on or after the international filing date
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
"O" document referring to an oral disclosure, use, exhibition or other means
"P" document published prior to the international filing date but later than the priority date claimed

- "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
"&" document member of the same patent family

Date of the actual completion of the international search

24 March 1995

Date of mailing of the international search report

30.03.95

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+ 31-70) 340-2040, Tx. 31 651 epo nl,
Fax (+ 31-70) 340-3016

Authorized officer

Nguyen Xuan Hiep, C

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 94/13551

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US-A-4277836	07-07-81	JP-C- 1238309	31-10-84
		JP-A- 52103925	31-08-77
		JP-B- 59011980	19-03-84
EP-A-0217479	08-04-87	JP-C- 1349993	28-11-86
		JP-A- 57121746	29-07-82
		JP-B- 61013268	12-04-86
		DE-A- 3278649	14-07-88
		EP-A, B 0057096	04-08-82
		US-A- 4575796	11-03-86
US-A-4028663	07-06-77	NONE	